content addressable storage; store metadata associated with the file in a storage mechanism for storing metadata for content addressable storage; receive a query from a requester for content; search the metadata storage mechanism for content related to the received query; and when the metadata associated with the file is indicated by the query: retrieve the file stored in the content addressable storage; and send the retrieved file to the requester.

In yet another embodiment, a computer-implemented system for managing metadata in a content addressable storage system, comprising: a content addressable storage system configured to receive a file for storage, said file to be stored using content addressable storage; and store metadata associated with the file in a storage mechanism for storing metadata for content addressable storage; and an application server configured to: receive a first query from a requester for content; search locally for content related to the first query; send a second query, related to the first query, to the content addressable storage system; receive one or more files related to the second query from the content addressable storage system; and send to the requester the one or more files found locally based on the first query and the one or more files received from the content addressable storage system based on the second query.

These and other features and advantages of the invention will become apparent from the following description of embodiments. Neither this summary nor the following detailed description purports to define the invention. The invention is defined only by the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

These and other features will now be described with reference to the drawings summarized below. These drawings and the associated description are provided to illustrate specific embodiments, and not to limit the scope of the invention.

FIG. **1** illustrates a block diagram of a system for metadata management in a content addressable storage system.

FIG. **2** is a block diagram representing an exemplary process for storing metadata in one or more CAS servers.

FIG. **3** is a block diagram of an exemplary process for replicating objects and metadata.

FIG. **4** depicts a block diagram of an exemplary process for retrieving CAS content from a CAS server.

FIG. **5** depicts a block diagram of an exemplary process for retrieving CAS data via an application server.

DETAILED DESCRIPTION

In the following detailed description, references are made to the accompanying drawings that illustrate specific embodiments in which the invention may be practiced. Electrical, mechanical, programmatic and structural changes may be made to the embodiments without departing from the spirit and scope of the disclosure. The following detailed description is, therefore, not to be taken in a limiting sense and the scope of the disclosure is defined by the appended claims and their equivalents.

FIG. **1** illustrates a block diagram of a system for metadata management in a content addressable storage system. In FIG. **1**, numerous computers and systems are all interconnected via a network **101**. The network **101** may be the internet, an intranet, a dedicated wired network, one or more cables, a wireless network, or any other appropriate type of communication. The system may include one or more CAS servers **110**A-C. Each of these CAS servers may include or have thereto attached one or more storage systems **120**A-C. The storage systems **120**A-C may include one or more RAID storage systems, cloud storage, tape storage, optical disks, magnetic disks, and/or any other appropriate type of storage. There may be any number of CAS servers and each may have any number of storage systems **120**A-C. Two or more storage systems **120**A-C may reside on the same physical disk or storage, or each storage system **120**A-C may reside on one or more disks or other storage separate from all of the other storage systems **120**A-C. In some embodiments, content stored in a CAS server **101**A-C may be retrieved based on a GUID, as is known in the art, or based on a search, as discussed herein.

As noted herein, when a CAS server **110**A-C receives content to store, it may store the content and metadata to the storage system **120**A-C. The content may be stored in the format in which it is received, or in flat files, directories, databases, or the like. The metadata may be stored in any fashion, including in a database, flat file, directory of files, or the like. In some embodiments, the metadata may be stored in XML or other structured file as plain text and this plain text may be searchable. In some embodiments, the metadata is stored in a database, and this database may be searchable.

The CAS servers **110**A-C may be coupled via the network to one or more application servers **130**. The application server **130** may include or have thereto attached a storage system **150**. The storage system **130**, like the storage system **120**A-C, may include one or more one or more RAID storage systems, cloud storage, tape storage, optical disks, magnetic disks, and/or any other appropriate type of storage. In some embodiments, the application server **130** is used to receive one or more files, make a decision to store the file in CAS and to send the file and or metadata to a CAS server **110**A-C in order to store the CAS content and/or the metadata. In some embodiments, the application server **130** may be used by a user using a client computer, e.g., client system **140**A, in order to query for content. The user may submit a query for content to the application server **130** and the application server may attempt to respond to that query both by looking locally, including on its storage system **135** and by sending the query to one or more CAS servers **110**A-**110**C.

In some embodiments, one or more client systems **140**B-**140**C are coupled to the network **101** and may allow a user to query the CAS servers **110**A-**110**C via a client application directly without going through an application server **130**.

Storing the CAS Content and Metadata

FIG. **2** is a block diagram representing an exemplary process for storing metadata in one or more CAS servers. In step **210**, content is received. In some embodiments, receiving content may include receiving DICOM images, images, historical data, video or other content. The content may be received, e.g., at an application server **130**.

Content Addressed Storage (CAS) is a technique by which a unit of data stored on a storage system is accessed using an address that is derived from the content of the unit of data. As an example, the unit of data may be provided as an input to a hashing function which generates a hash value that is used as the content address for the unit of data. When a host computer sends a request to a content addressable storage system to retrieve a unit of data, the host provides the content address (e.g., hash value) of the unit of data. The storage system then determines, based on the content address, the physical location of the unit of data in the storage system, retrieves the unit of data from that location, and returns the unit of data to the host computer.

If, in step **220**, a determination is made to not store the content in CAS, then the application server **130** may store the content locally in a storage system **135**. The application